

Imate

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application No.: 10/659,181
Filed: September 11, 2003
Inventor:
Scott Douglas Kaiser

Title: Data Storage Analysis
Mechanism

§ Examiner: Unknown
§ Group/Art Unit: 2862
§ Atty. Dkt. No: 5760-13600/VRTS
§ 0412

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date indicated below.

Robert C. Kowert

Printed Name

[Signature]

Signature

April 1, 2004

Date

PETITION TO MAKE SPECIAL UNDER 37 C.F.R. § 1.102(d)

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

This is a petition to make special under 37 C.F.R. § 1.102(d) and pursuant to MPEP 708.02(VIII). Please charge the petition fee under 37 C.F.R. § 1.17(h) to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5760-13600/RCK.

A pre-examination search was performed by a search firm. The search firm reported that the field of search was class 702, subclass 182; class 709, subclasses 223 and 224; class 710, subclass 1; and class 711, subclasses 100 and 170; and also included extensive computer searching in various U.S. and foreign patent related computer databases including the U.S. Patent and Trademark Office computer database. The search results returned to the undersigned by the search firm are nine U.S. patents or published patent application references, copies of which are included herewith. The

references are discussed in detail below, and the patentability of the claimed subject matter is pointed out.

Detailed Discussion of References

U.S. Patent 6,269,410 to Spasojevic

Spasojevic is generally concerned with using system traces to characterize workloads in a data storage system. Particularly, Spasojevic teaches “I/O activity of a data storage system is characterized by collecting system traces generated during the I/O activity, grouping records in the system traces according to stores, identifying I/O activity in streams corresponding to the stores, and processing the groups of records to characterize I/O activity for the streams. The stores represent units of storage such as logical drives, single data storage devices and groups of data storage devices in the data storage system. Characterizations of the I/O activity may be used to re-allocate data across the data storage system.” (Spasojevic, abstract).

Spasojevic further teaches: “an enterprise system including a data storage system and a machine that makes use of system traces to characterize I/O activity of different streams in the data storage system. The data storage system includes a range of data storage devices. Among the characterizations generated from the system traces is an overlap in I/O activity across different streams. The I/O activity characterizations may be used to allocate data more efficiently across the data storage system.” (Spasojevic, column 2, lines 49-57). Spasojevic further teaches: “During post-processing, records in the system traces are grouped according to stores, I/O activity in streams is identified and groups of records are processed to generate attributes that characterize I/O activity of the different streams. Based on these I/O activity characterizations, data can later be re-allocated across the data storage system 12.” (Spasojevic, column 4, lines 51-57). Spasojevic further teaches: “Characterizing the I/O activity in this way can significantly improve the worst-case analysis of the performance of the data storage system 12. The requestRate, requestSize and runCount attributes may be used to characterize needed throughput in the data storage system 12. In addition, the overlap Fraction attribute might be used to determine better Quality of Service (“QoS”) guarantees regarding response

time, than what is feasible based just on requestRate, requestSize and runCount attributes.” (Spasojevic, column 5, lines 3-11).

Spasojevic does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Spasojevic for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Spasojevic for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Spasojevic for at least the above reasons.

U.S. Patent 6,606,585 to Borowsky et al. (“Borowsky”)

Borowsky is generally concerned with acceptability testing for capacity planning of data storage system. Particularly, Borowsky teaches that “Data storage devices of an enterprise system are tested to determine whether the enterprise system is optimally configured. Each data storage device is tested to determine whether it can satisfy a performance requirement for an assigned group of n workloads. A group of n inequalities are generated, and only up to n of the inequalities may be evaluated to determine whether the device satisfies the performance requirement for the assigned group of workloads.

The inequalities are based on a phased, correlated model of I/O activity.” (Borowsky, abstract).

Borowsky further teaches: “a method and apparatus for testing whether performance requirements are satisfied if a given group of workloads is assigned to a data storage device. The test is based on a phased, correlated model of I/O activity in the group of workloads assigned to the device.” (Borowsky, column 3, lines 31-35). Borowsky further teaches: “A probability of overlap $p_{sub,ji}$ can be modeled, measured or estimated to indicate the degree of simultaneous activity between the $i_{sup,th}$ and $j_{sup,th}$ streams, where $j_{lto req, n}$ (e.g., the first and $n_{sup,th}$ streams S_1 and S_n). The probability $p_{sub,ji}$ represents the probability that the $i_{sup,th}$ workload S_i is already ON when $j_{sup,th}$ workload S_j comes ON. The probability $p_{sub,ji} = 0$ corresponds to an $i_{sup,th}$ workload S_i that is never ON when a $j_{sup,th}$ workload S_j comes ON. The probability $p_{sub,ji} = 1.0$ corresponds to an $i_{sup,th}$ workload S_i that is always ON when a $j_{sup,th}$ workload S_j comes ON. These probabilities $p_{sub,ji}$ can be used in a test that determines whether the data storage device 12 can satisfy one or more performance requirements for an assigned group of workloads S_1 to S_n . The test is based on a phased, correlated model of I/O activity in the group of workloads S_1 to S_n assigned to the data storage device 12. Each workload S_1 to S_n is assumed to exhibit phase behavior (i.e., turn ON and OFF). The performance of the data storage device 12 may be measured by attributes such as stability, capacity, speed and Quality-of-Service guarantees.” (Borowsky, column 4, lines 30-49).

Borowsky does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Borowsky for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage

availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Borowsky for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Borowsky for at least the above reasons.

U.S. Patent 6,636,905 to McNamer et al. (“McNamer”)

McNamer is generally concerned with analyzing input/output performance of a data processing system. Particularly, McNamer teaches a “method for analyzing input/output performance of a data processing system. The method comprises providing records of input/output operations performed by the data processing system. Each record has statistics related to a respective input/output operation for a respective process. One or more workload classes are specified for accumulating statistics, and the statistics from the input/output records are accumulated by workload class. The accumulated statistics are reported by workload class which supports analysis of the input/output behavior for different workload classes.” (McNamer, abstract).

McNamer further teaches: “In an example embodiment, a file of input/output operations that spans a selected period of time is provided as input. Using workload classes (e.g., process types) that are user-specified, statistics are accumulated by workload class for the record of input/output operations. The workload classes of the process identifiers in the record of input/output operations are obtained from a system table having associations of process identifiers to workload classes. The input/output statistics are then reported by workload class. Thus, a user can see the input/output statistics for a present system configuration, and the statistics can be further provided as

input to a modeling tool to see how a new system configuration would impact performance of the input/output operations.” (McNamer, column 3, lines 31-35).

McNamer further teaches: “I/O analyzer 140 accumulates statistics according to workload classes specified by category selections file 130.” (McNamer, column 4, lines 1-2). “I/O analyzer 140 produces report file 170 indicating various accumulated I/O statistics.” (McNamer, column 4, lines 22-23). “The data from report file 170 and statistics file 180 can be provided as input to I/O modeling tool 190. I/O modeling tool projects performance level increases for the selected workload classes if the system configuration is modified in a predetermined manner...Using predetermined scaling factors for the hardware to be added to the configuration, a projected response time can be determined and presented.” (McNamer, column 5, lines 27-36).

McNamer does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in McNamer for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in McNamer for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in McNamer for at least the above reasons.

U.S. Patent Application Publication 2003/0023713 to Slater et al. ("Slater")

Slater is generally concerned with a monitoring appliance for data storage arrays, and monitoring the usage of such arrays. Particularly, Slater teaches a "monitoring appliance for a data storage array used by plural hosts to store data responds to stored metadata to interrogate the data storage array at intervals to establish the amount of usage of the data storage array. Each host can use the file system(s) and/or database(s) of its choice in portions of the data storage array allocated to it. The monitoring appliance has basic knowledge of all file systems/databases used by the hosts, and the metadata structure of those file systems/databases." (Slater, abstract).

Slater further teaches: "The appliance 12 interrogates the disc array 10, and...determines the form of, structures of, and capacities of, the file systems/databases being stored on the host allocated areas of the disc array 10. Appliance 12 stores the metadata describing this information within a storage device in its operating environment. By doing the interrogation, determining and storing steps, appliance 12 self configures an internal database of the host allocated areas of the disc array 10 and the related metadata of the file systems/databases and their initial levels of usage." (Slater, page 2, paragraphs 31-32). Slater further teaches: "After appliance 12 has stored the describing metadata, the appliance 12 re-interrogates disc array 10 from time to time preferably at regular time intervals T, e.g. every few minutes, to establish the then current usage of the file systems/databases by obtaining the current values of the metadata, which are added to the internal database of the appliance. Comparisons are then made between the original state of the file systems/databases and their state at any later time that the metadata is obtained." (Slater, page 2, paragraph 33). Slater further teaches: "In an alternative embodiment, the appliance 12 dual mounts the file systems and/or databases which are on the disc array 10. Appliance 12 only reads the file system and/or databases on array 10 in such a manner that the functioning of the data storage within the disc array 10 cannot generally be disrupted. At regular intervals thereafter appliance 12 compares the current state of the file systems and/or databases with the initial state thereof, by comparing the file systems and/or databases at the two times." (Slater, page 2, paragraph 34). Slater

further teaches: “With either manner of operation of the appliance 12, updates concerning the usage of the disc array 10 are posted to management station 16 via the management interface. The posting may conveniently use a simple kind of web publishing, such as an HTML web page, although any appropriate form that contains basic capacity usage information can be used. Appliance 12 posts such reports from time to time, e.g., at predetermined intervals, or as and when demanded by the management station.” (Slater, page 3, paragraph 35).

Slater does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Slater for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Slater for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Slater for at least the above reasons.

U.S. Patent Application Publication 2003/0061546 to Collins et al. (“Collins”)

Collins is generally concerned with monitoring the performance of a storage device. Particularly, Collins teaches an “apparatus and method for monitoring performance of a storage device. Preferably, the apparatus is embodied in computer readable program code. The apparatus and method may intercept communications

between a computer system and the storage device, wherein the communications are compared to a threshold value and/or a predicted failure of the storage device. The apparatus and method may also respond to a decline in the performance of the storage device based on the analyzed communications, and preferably prior to the predicted failure thereof.” (Collins, abstract).

Collins further teaches: “The computer readable program code may comprise: program code for intercepting communications between the computer system 100 and the storage device 150; program code for analyzing the intercepted communications, wherein the intercepted communications are compared to a predicted failure of the storage device 150; and program code for responding to a decline in the performance of the storage device 150 prior to the predicted failure thereof.”

Collins further teaches: “an information log 200 for monitoring the performance of the storage device 150” which may include information for “correcting the access time for overhead...determining a threshold for responding to a decline in performance...responding to a decline in the performance of the storage device 150, prior to failure thereof...monitoring the performance of the storage device 150 and responding to a decline in performance thereof prior to a failure of the storage device 150 [and] monitoring the performance of the storage device 150 and responding to a decline in performance thereof by defragmenting all or a portion of the storage device 150.” (Collins, pages 2-3, paragraph 25).

Collins further teaches: “the storage device 150 is monitored for “hidden” or “masked” signs of declining performance...a recoverable failure may be identified based on an analysis of the failures reported from the storage device to the computer system 100. That is, the filter driver 130 may intercept reported errors or failures 170. In addition, the filter driver 130 and/or other suitable program code may also intercept other communications, such as, the location of attempted access on the storage device 150, amount of data, type of access, duration of access, etc. The information from the intercepted communications may be written to a storage database (e.g., information log

200). Suitable program code may also be provided for analyzing the intercepted communications. For example, the program code for analyzing the intercepted communications may find an increasing number of failures for the storage device 150 and/or a particular area thereof. When the number of failures exceeds a threshold, program code for responding to the decline in performance of the storage device 150 may warn the user of a potential or pending problem with the storage device 150.” (Collins, page 3, paragraphs 26-27).

Collins further teaches: “the intercepted communications may be compared to a known or predicted failure 375 of the storage device 150. That is, based on past performance of comparable storage devices, it may be known that the storage device 150 may fail entirely when the storage device 150 experiences a number of recoverable failures...Alternately, or in addition to, the failure 375 may be derived or predicted to fail entirely when the storage device 150 experiences a number of recoverable failures...based on statistical analysis of the logged information (e.g., curve fit 355). A performance threshold 370 may be determined based on the known or predicted failure 375. Thus, when the analysis of the intercepted communications indicate that the performance of the storage device 150 is approximately at the performance threshold 370, a response may be initiated to prevent loss of the data on the storage device 150...As such, a response may be initiated prior to failure of the storage device 150 to prevent, or reduce the risk of, losing the data on the storage device 150 due to a failure thereof.” (Collins, page 3, paragraphs 30-31).

Collins does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Collins for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level

under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Collins for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Collins for at least the above reasons.

U.S. Patent Application Publication 2003/0145186 to Szendy et al. (“Szendy”)

Szendy is generally concerned with measuring and optimizing spatial segmentation of electronic storage workloads. Particularly, Szendy teaches a “method and apparatus...for measuring and optimizing the orientation of data access of an electronic storage device according to data access characteristics. Monitoring storage access activity in an area of storage space is performed to gather data pertaining to one or more storage access characteristics. Measuring is performed of the characteristics of the storage access activity of at least two individual portions of the storage space. The portions are then combined in a manner to more judiciously utilize storage space. Depending on their homogeneity of access characteristics, the individual portions may then be left alone, merged with other similar portions, or further subdivided into sub-portions, which may be further merged, divided or left alone. At each merger or division determinations can then be made of whether the characteristics of storage access activity of one individual portion or sub-portion are similar to that of another portion according to predetermined criteria.” (Szendy, abstract).

Szendy further teaches: “a method and apparatus for measuring and optimizing the data access activity of electronic storage devices. Generally, the method includes an analysis of the I/O activity of storage space. In a disk drive storage device for example, storage space may be analyzed according to the I/O activity in its disk space. The space may then be divided or merged from a spatial point of view. The divisions and mergers of

portions of the storage area are done according to access characteristics of individual portions of the storage space, including the I/O activity. The I/O activity may include the number of I/Os in a particular portion, the rate of I/Os over a period of time, the ratio of reads to writes, the locality of I/O activity, access patterns, burstiness of the access, and other characteristics. These characteristics may be measured and used to optimize the use of the storage medium and to aid device and system designers in laying out designs.” (Szendy, page 2, paragraph 14). Szendy further teaches: “In one embodiment, the information may be used by a device, a personal computer for example, to dynamically monitor and manage data access of storage devices. If neighboring portions are found to have similar access characteristics, they may be merged from a spatial perspective. If they are found to be different, they may be left alone.” (Szendy, page 2, paragraph 15).

Szendy further teaches: “In one embodiment of the invention, an online tool may be embedded in a device for managing storage activity. In such an embodiment, the tool would be able to monitor storage space access patterns as they occur. Since these patterns typically change over time, an embedded tool that incorporates the invention is able to monitor storage accesses as needed in order to more efficiently manage the use of storage space. In one embodiment, the tool can take periodic samples of access characteristics to determine whether the storage space needs to be reapportioned according to its use. In another embodiment, every access operation may be monitored in order to get a more complete understanding of the access characteristics.” (Szendy, page 2, paragraph 19).

Szendy does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Szendy for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16

recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Szendy for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Szendy for at least the above reasons.

U.S. Patent Application Publication 2004/0015908 to Giel et al. (“Giel”)

Giel is generally concerned with analyzing the configuration information of an enterprise based on expert knowledge. Particularly, Giel teaches a “computerized apparatus and method [that] generates information descriptive of issues arising in a monitored set of nodes. The method comprises supplying one or more analyzer programs with node-specific information derived from one or more nodes. Each analyzer program is then caused to analyze at least some of the information provided from each node to detect the presence of one or more issues. When an issue is detected, the analyzer program is caused to generate issue identification information which is augmented with information identifying the node from which the node-specific information was derived. The augmented information for all the issues detected is presented in combined form as an issues database suitable for later use in report generation.” (Giel, abstract).

Giel further teaches: “an automated method 100 for repeatedly analyzing the configuration of an enterprise. The method 100 may generally be broken into three stages: collection, analysis, and reporting... enterprise configuration information is collected from field nodes. The collection step utilizes a set of collectors 104 to gather the desired configuration information...The collectors 104 gather desired configuration information and store it as files in the tracker database 106... configuration information is analyzed by an analyzer harness 806 (FIG. 2) to identify issues within the field nodes. As part of this process, the analyzers 110 are used to investigate particular issues associated with the field nodes... the analyzer harness 806 executes the desired analyzer 110 with

the configuration information stored in the tracker database 106 and generates a report 2126 (FIG. 21) in XML format. The report 2126 identifies issues relating to the field nodes. This issue identifying report 2126 is then stored as a file in an issue database 112... At this stage, the report 2126 generated by step 108 may be used to generate a full report along with text descriptions of the issues, as desired by an auditor.” (Giel, page 4, paragraphs 72-77).

Giel further teaches: “the issues stored in the issues database 112 may be used in another way to review the overall performance of the enterprise. To this end, in step 114 the issues are analyzed using rules written by the experts, and a report is generated as desired by the auditor...The reports may be presented in step 118 to enterprise management, technical management, the field engineering team, and to a workflow system or healer system (self-healing technology).” (Giel, page 4, paragraph 78).

Giel does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Giel for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Giel for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Giel for at least the above reasons.

U.S. Patent Application Publication 2004/0024921 to Peake, JR. et al. ("Peake")

Peake is generally concerned with analyzing input/output activity on local attached storage. Particularly, Peake teaches a "system, method and computer program product for analyzing file I/O activity on local attached storage devices within a computer network is provided. In an embodiment, a software agent executes on one or more servers within the network, and monitors the I/O activity on the network's local attached storage (e.g., SAN, NAS, and IDE and SCSI disks). A management interface is also provided for monitoring I/O activity-related data and for receiving reports on such I/O activity. In an embodiment, collected I/O-related data and any predefined I/O metrics are stored in a central repository (e.g., a relational database). The system, method and computer program product provide accurate metrics to assists system administrators in deciding, justifying and validating resource purchases for and allocations within the network." (Peake, abstract).

Peake further teaches: "The present invention provides a system, method and computer program product for analyzing I/O activity. In an embodiment, a software agent is provided that executes on one or more servers within a network, where each server is running the Microsoft.RTM. Windows 2000.TM. or XP.TM., IBM.RTM. AIX.TM. or Sun.RTM. Solaris.TM. operating system. The software agent monitors the I/O activity on the network's local attached storage device (i.e., one or more disk volumes). A management interface is also provided for monitoring I/O activity-related data and for reporting purposes. In an embodiment, collected I/O-related data and any predefined I/O metrics are stored in a central repository. In one embodiment, the central repository is a relational database (e.g., Oracle9i.TM. or Microsoft.RTM. SQL Server.TM. database) residing on a database sever which is separate from the application server where the I/O analyzer agent is executing." (Peake, page 2, paragraphs 24-26).

Peake further teaches: "In step 220, process 200 determines if any I/O metric(s) have fallen outside any threshold(s) set by the administrator (e.g., in step 204)." (Peake, page 6, paragraph 80). Peake further teaches: "In step 222, process 200 would perform

any threshold actions set (e.g., defined in step 204) for each of the thresholds determined to have been exceeded in step 220.” (Peake, page 7, paragraph 82). Peake further teaches: “The purpose of alerting is to send the user (e.g., system administrator) a message whenever I/O performance of a managed object falls outside of normal limits (i.e., thresholds).” (Peake, page 7, paragraph 84).

Peake further teaches: “In an embodiment, the user may perform data analysis to identify applications that are causing I/O bottlenecks and obtain information to allow optimization of such applications so that these bottlenecks can be eliminated or reduced. This can be done both by allowing interactive data analysis and with predefined reports.” (Peake, page 8, paragraph 88).

Peake does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Peake for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Peake for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Peake for at least the above reasons.

U.S. Patent Application Publication 2004/0025162 to Fisk

Fisk is generally concerned with managing application workloads and data storage resources. Particularly, Fisk teaches “methods and associated systems for managing application workloads and data storage resources. Techniques are disclosed for determining the I/O capacity of a data storage resource for a given workload and allocating resources according to administrator requirements. The invention may be implemented as a transparent layer between the application and the data storage resource, for example, in the file system.” (Fisk, abstract).

Fisk further teaches: “In accordance with one embodiment of the invention, the data storage manager 118 estimates the response time of a workload on a data storage resource servicing a complex workload. A system administrator may use this information to configure the system to provide a desired level of performance. In accordance with one embodiment of the invention, the data storage manager allows the administrator to define a desired system performance and then estimate, based on system I/O performance, an appropriate mapping of data storage resources to application workload to achieve the desired system performance. For example, a desired level of performance may be achieved by spreading the workload of an application across several data storage devices. The teachings herein may be used to determine, for a given data storage resource and complex workload, the number of data storage devices across which a particular application workload should be spread.” (Fisk, page 3, paragraphs 53-54).

Fisk further teaches: “operations for analyzing response times for complex workloads associated with a data storage resource...the data storage manager 118 analyzes the data storage resource to determine the response characteristics (e.g., response times) of the data storage resource for various workloads...This analysis also may take into account the load level in the system. That is, data may be collected to determine the response characteristics as the number of concurrent requests varies...the data storage manager 118 determines the load level (e.g., average number of concurrent operations) for the analysis. Typically, this is an empirical measurement of the load level in the system. Alternatively, the load level may be selected if, for example, the system is

being configured to provide a certain level of performance at that load level...the data storage manager 118 determines the complex probability distribution of the application workload...the data storage manager 118 calculates the I/O capacity of the data storage resource for a given workload... the data storage manager 118 calculates the utilization level of the data storage resource 116. The utilization level is a measure of percent of I/O capacity that is being used.” (Fisk, page 4, paragraphs 55-61).

Fisk further teaches: “operations for configuring system resources in response to measured response times and/or utilization levels...the administrator defines a desired utilization level and inputs this into the data storage manager 118. Next, the data storage manager 118 determines the load level in the system (block 216)...the data storage manager 118 determines the I/O capacity of the data storage resource 116...the data storage manager 118 calculates the desired parameter. In this example, the parameter is the stripe width necessary to provide a given utilization level. It is possible, however, that a given system may not provide the desired level of performance. In this case, the data storage manager 118 may notify the administrator that the system S may need to be reconfigured. For example, new data storage resources may be added or the data storage resources 116 may be replaced with different data storage resources. In one embodiment, the data storage manager 118 may automatically reconfigure the system S.” (Fisk, page 4, paragraphs 62-65).

Fisk further teaches: “In an alternate embodiment, a process similar to that of blocks 212-220 may be used to solve for the load level in the system. For example, given a desired utilization level and data storage resource 116, the data storage manager 118 may determine the number of concurrent streams necessary to support the desired performance.” (Fisk, page 4, paragraph 66).

Fisk does not teach or suggest the combinations of features recited in each of claims 1-33 of the captioned application. For example, claim 1 recites a combination of features including: “program instructions...executable by the processor to: determine a storage requirement for the storage system to meet a given storage availability risk level

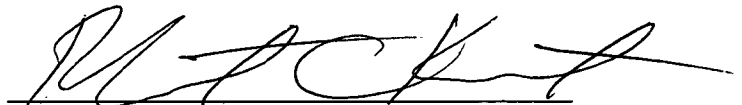
under one or more conditions indicated by the storage demand data.” Claims 2-14 depend from claim 1 and thus are not taught or suggested in Fisk for at least the above reasons. Claim 15 recites a combination of features including “means for determining a storage requirement for the storage system to meet a given storage availability risk level under the one or more conditions indicated by the storage demand data.” Claim 16 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 17-24 depend from claim 16 and thus are not taught or suggested in Fisk for at least the above reasons. Claim 25 recites a combination of features including “determining a storage requirement for the storage system to meet a given storage availability risk level under one or more conditions indicated by the storage demand data.” Claims 26-33 depend from claim 25 and thus are not taught or suggested in Fisk for at least the above reasons.

CONCLUSION

Applicant requests that the application be granted "special" status and that handling of the application be accelerated correspondingly within the Office.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/13600/RCK.

Respectfully submitted,

A handwritten signature in dark ink, appearing to read 'RCK', is written over a horizontal line.

Robert C. Kowert
Reg. No. 39,255
AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8850

Date: April 1, 2004